Life Is On | Schneider Electric

Direct Liquid Cooling System Challenges in Data Centers

White Paper 210

Version 1

Data Center Research & Strategy

by Maria A. Torres Arango
Victor Avelar
Stefano Lena
David McGlocklin

Executive summary

Direct liquid cooling (DLC) has become the IT industry's preferred cooling method for extreme chip power densities. DLC systems are relatively new to data centers and pose some challenges to data center professionals. This paper describes the eight most common liquid cooling system challenges related to specification, installation, and operation. The paper then provides guidance for each challenge to help with liquid cooling deployments.

RATE THIS PAPER ★★★★★

Introduction

Artificial intelligence (AI) and high performance computing (HPC) workloads depend on high-performance chips (typically graphics processing units or GPUs) to perform mathematical calculations. As the computing performance for these chips intensifies, more heat must be removed. Server vendors are switching from air-cooling to liquid-cooling to deal with this increased heat density. Direct liquid cooling (DLC), also known as direct-to-chip, has become the de facto method for cooling these chips.[1] However, the data center industry is still learning how to implement this technology at scale and is in the early phases of developing standards around it. As result, there are some key challenges related to DLC systems in data centers.

In this paper, we focus on the challenges associated with DLC applications involving around 500 kW or more and 10 or more IT racks. These rough thresholds indicate where a specific type of DLC architecture becomes applicable, as we will explore in detail later.

Challenges

Table 1 lists the eight challenges discussed in this paper. They are organized by life phases: Specification, Installation, and Operation. They are also ranked within each life phase by importance, highlighting the most critical ones first. Several of the challenges stem from a lack of industry standards. Not only do standards serve manufacturers (e.g., testing, ratings), they also help define how to specify, install, and operate these key technologies.

We describe each challenge and provide guidance on how it should be addressed. Note, the challenges in this paper are not an exhaustive list, but ones that are top-of-mind today. Other challenges will arise as adoption grows and technology continues to evolve. **For convenience, each challenge in the table is hyperlinked, so you can easily navigate to that section.** Also, every page has a "home" symbol on the upper right which returns you to Table 1.

Table 1

Direct liquid cooling system challenges categorized by specification, installation, and operation

	Direct liquid cooling system challenges
Specification	
1	Material incompatibility between CDU and connected components increases risk of server damage
2	Opposing liquid and air-cooling requirements forces a tradeoff between energy and capital costs
3	Direct coupling between servers and DLC infrastructure complicates specification
4	Lack of CDU system efficiency standards complicates comparisons
5	Provisioning IT space for unknown liquid-cooled IT risks higher costs from stranded capacity
Installation	
6	Complexity of preventing TCS contamination increases the risk of server damage
Operation	
7	Lack of clarity between cooling and server vendor warranties creates undue stress
8	Slow DLC system response to GPU power transients poses a risk of GPUs overheating

[1] In our executive brief, *Optimizing AI, the Critical Role of Liquid Cooling*, we discuss the growing need for liquid cooling in data centers and the importance of careful planning and a strong ecosystem.

Liquid cooling system and CDU fundamentals

In this section, we describe DLC architectures and coolant distribution units (CDUs), to provide a foundation for the remainder of the paper. (If you're familiar with these concepts, skip to the next section.) As we state in White Paper 133, *Navigating Liquid Cooling Architectures for AI Workloads*, a DLC architecture can be described by three elements (Figure 1):

1. Heat capture within the server
2. CDU type and coolant
3. Method of rejecting heat to the outdoors

Figure 1

Simplified view of a liquid cooling architecture in a data center

Source: ASHRAE, Water-Cooled Servers: Common Designs, Components, and Processes, page 10

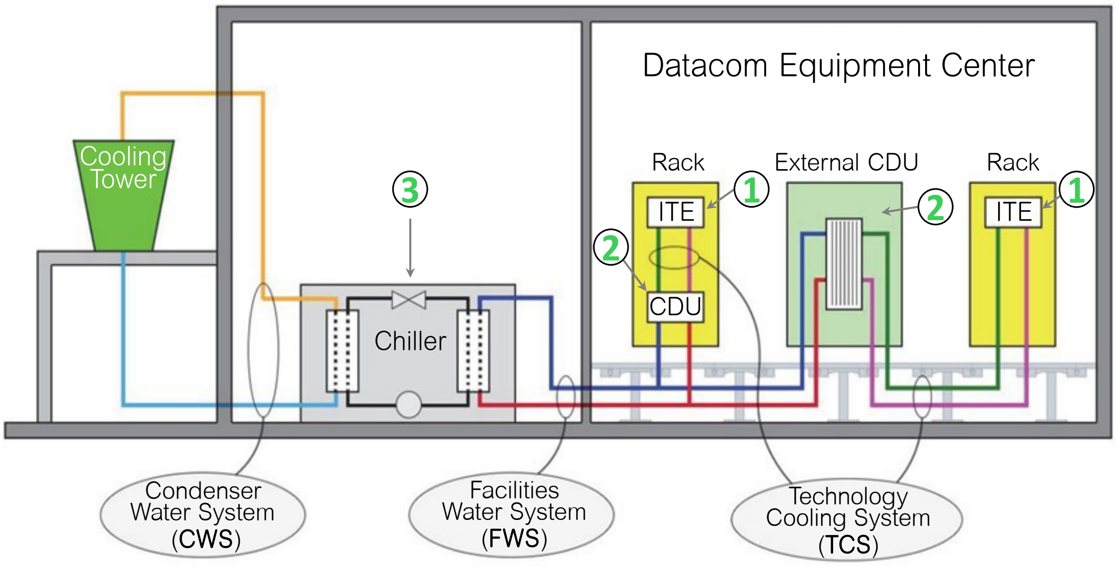

Heat capture within the server – This is achieved with either cold plates directly attached to components (i.e., GPUs) or with immersion cooling (not discussed in this paper since DLC is the prevalent means in the industry). Coolant flowing through the cold plates transports the heat energy out of the server and into the CDU.

CDU type and coolant – Liquid-to-liquid (L-L) and liquid-to-air (L-A) CDUs are the most prevalent, but there are four others. L-A CDUs are the easiest to add to conventional air-cooled data centers. This is because the heat removed from the IT equipment is released to the air in the white space for the air-based cooling system to manage. However, this option is often more cost effective for smaller liquid cooling deployments of 1 to 10 racks.

In this paper, we focus on larger deployments using L-L CDUs. With L-L CDUs, pumps circulate coolant through the IT equipment. The coolant absorbs the IT heat and transports it to the CDU heat exchanger. The heat exchanger transfers the heat from the technology cooling system (TCS) to the facility water system (FWS). Hence the name "liquid-to-liquid". L-L CDUs can support large capacity deployments (>500 kW). The TCS is the dedicated cooling system that removes the heat from the IT equipment and transfers it to the facility cooling system. It is inclusive of the CDU(s), IT equipment, piping, manifolds, valves, connectors, coolant quality monitoring & treatment, and controls.

A CDU provides five key functions (i.e., temperature control, flow control, pressure control, coolant filtration and monitoring, heat exchange and isolation). The isolation function is critical because the IT coolant flows through cold plates in the IT equipment. FWS water carries larger particles than TCS coolant. These cold plates have very small channels that can easily become clogged, placing the chip at risk of overheating and damage, thus requiring very stringent control of the coolant composition. This is why the requirements for the FWS and TCS water[2] are different. Figure 2 illustrates a concept of an AI or HPC pod deployment with the CDUs in the server corridor.

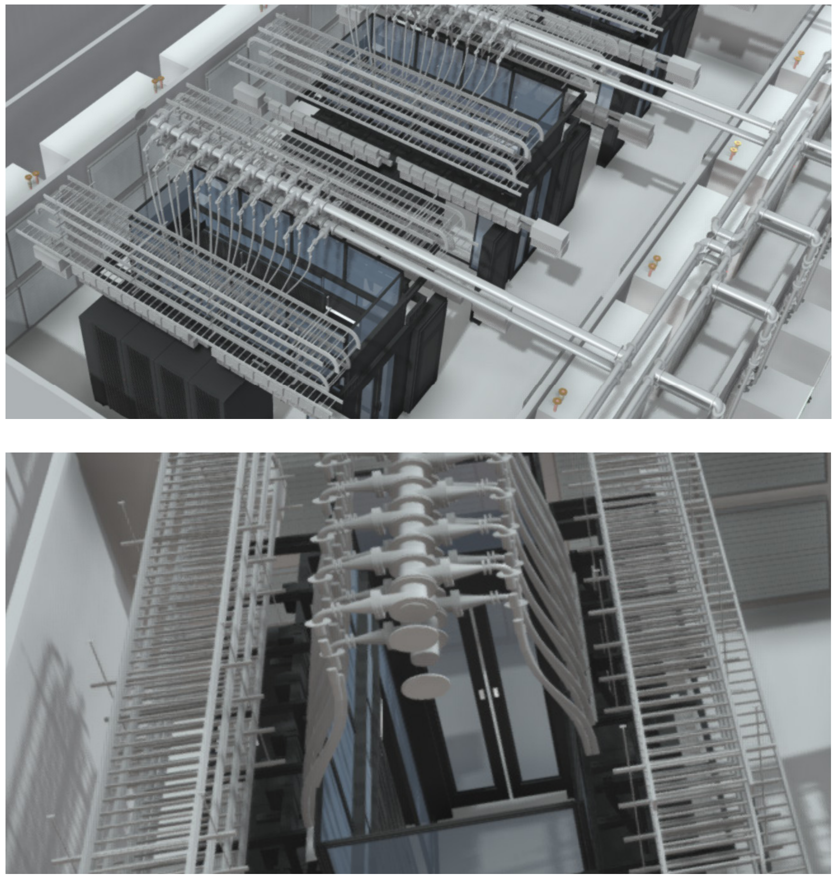

Figure 2

Illustration of a liquid-cooled pod showing coolant distribution from the CDUs in the service corridor to the pod.

Method of rejecting heat to the outdoors – There are several data center heat rejection systems in use today. Chilled water is among the most common in larger data centers. This includes air-cooled or water-cooled chillers with pumps that circulate water through the facility side of the CDU and through traditional computer room air handlers (CRAH). Another common heat rejection system uses a dry cooler. Sometimes dry coolers use water spray to lower the facility water temperature (i.e., adiabatic cooling).

[2] The term "water" when used in the context of data center cooling systems (FWS &, TCS), is a treated water variant including a mixture of water, propylene glycol, and other additives to prevent freezing and biological growth. ASHRAE, Water-cooled servers, Table 1, details the water quality guidelines.

Specification

A DLC system must be specified based on liquid-cooled IT requirements, but also considering your specific data center conditions. The combination of liquid-cooled IT with traditional heat rejection systems does present challenges. The first step, critical for the successful implementation of DLC systems, is specification. In this section we discuss its challenges.

1. Material incompatibility between CDU and connected components increases risk of server damage

When two or more metals with different electrochemical potentials are connected in the presence of a fluid, it sets off a process of galvanic corrosion. These metals are said to be *incompatible* when the difference becomes large enough to cause issues such as material degradation and redeposition. Galvanic corrosion is the exchange of metal ions from one metal to the other through the fluid.

The fluid in this case is the coolant circulated through the server's cold plates. All materials in the CDU, and connectors, seals, piping, valves, and the server's cold plates, in contact with the coolant are known as wetted materials. If these materials are incompatible, corrosion is triggered, eroding materials at different sites of the TCS loop. This creates debris within the coolant, representing a serious risk of damage to the IT equipment. Debris can clog the tiny channels in the cold plates and abrade their surfaces. Corrosion also increases the risk of leaks. Continuous damage also occurs in other components as the suspended debris wears against all surfaces (for example the CDU heat exchanger). Also note that the higher coolant temperatures used in these servers accelerate the corrosion.

GUIDANCE:
The selection of materials for the TCS loop is of utmost importance to prevent the issues described above. The key to material selection is to minimize the difference in the electrochemical potential between them. This can be done by using the anodic index.

The anodic index (measured in volts) for a material is the difference between its electrochemical potential in relation to that of gold. Lower (close to 0 V) anodic index means high material nobility, i.e., the material will better resist corrosion. Conversely, large anodic index values indicate that the material is less stable, i.e., the material is more prone to corrosion. Figure 3 illustrates the anodic index[3] of common materials. Acceptable metals listed by the Open Compute Project (OCP)[4] are indicated by green dots.

When using pairs of dissimilar materials in a system like the TCS loop, the difference in anodic index between the materials used should be minimized. This includes all wetted surfaces. For example, if the IT cold plates are made of copper (anodic index of 0.35 V), and the CDU heat exchanger is made of brass (anodic index of 0.35 V), avoid piping made of aluminum or galvanized steel, because their anodic index is too high (0.75-1.2 V), increasing risk of corrosion.

[3] Roberge, P., *Handbook of corrosion engineering*, 2nd edition. McGraw-Hill. 2012

[4] OCP, *Guidelines for using propylene glycol-based heat transfer fluids in single-phase cold plate-based liquid cooled racks*, Table 2., page 11.

Also, to attach piping sections, consider using the same piping material for connectors, or welding[5]. This keeps the list of materials used in the TCS loop as short as possible, which will reduce the risk of galvanic corrosion. Similar to OCP, the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) also provides a list of common wetted materials[6] found in TCS loops.

Figure 3

Anodic index of common materials. Materials with lower index values are more resistant to corrosion. High compatibility against galvanic corrosion exists for material pairs with similar anodic index. Acceptable metals for TCS wetted surfaces listed by OCP are indicated by green dots.

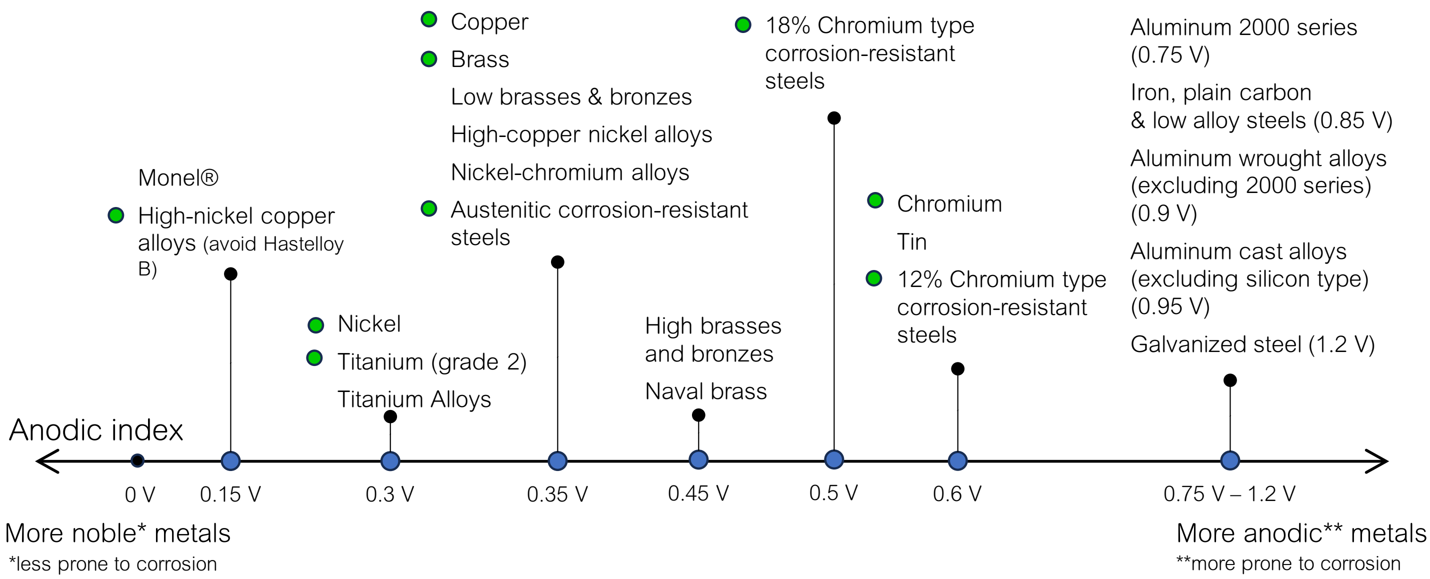

We recommend the following best practices when specifying materials within the TCS loop, to avoid the challenges presented earlier:

- Refer to compatible materials lists for TCS wetted materials (including the CDU) provided by IT equipment manufacturers. These guidelines should be strictly followed to prevent technical issues and maintain warranty coverage. If uncertain, speak to the IT vendor for clarification.
- Use the same or similar materials throughout the entire TCS loop when possible, to minimize incompatibility. Consider designing separate TCS loops if IT equipment cold plates materials vary (for example, IT equipment from multiple vendors).
- Maintain a registry of all materials used in the TCS loop design, including materials used in the CDUs and IT cold plates. Use this registry as reference for planning installation and maintenance tasks; and update periodically.
- Avoid aluminum as it has a "high" anodic index, or if unavoidable, make sure appropriate corrosion inhibitors are added to the coolant. Validate the materials selection with both the IT equipment and CDU vendors.
- Use your cooling solution provider's expertise to guide the design of your TCS loop if in-house expertise doesn't exist. The vendor's expertise should facilitate a more reliable deployment of liquid cooling in your data center and avoid potential warranty issues from choosing incompatible materials.

---

[5] High temperature joining processes such as welding require passivation to restore the anticorrosion properties of the materials, that may have been degraded from their exposure to high temperatures. Source: *Water-cooled servers*, 2019, ASHRAE

[6] ASHRAE, *Water-cooled servers*, Table 2, page 28.

2. Opposing liquid and air-cooling requirements forces a tradeoff between energy and capital costs

Liquid-cooled IT brings energy savings that aren't fully attainable when a single chiller plant supports both liquid-cooled and air-cooled IT loads. Compared to air, water is more than 23 times better at conducting heat (thermal conductivity) and can hold over 3,000 times more heat by volume (thermal capacity).[7] Putting these facts into relatable experiences will better explain the challenge this presents.

- Let's start with "water is more than 23 times better at conducting heat". If you accidently burn your finger on a hot stove, would you rather place your finger in 0°C (32°F) air or 10°C (50°F) water? Water provides more relief from the pain than air because it is better at *conducting* heat away from your finger, despite it being at a higher temperature than the air.

  Takeaway – Water cooling can effectively transfer heat from a processor even when the water temperature exceeds typical air-cooling temperature.
- Now on to "water can hold over 3,000 times more heat by volume". Imagine a cubic meter of asphalt and a cubic meter of water at room temperature. Now assume the summer sun shines on both for 6 hours. Which would you rather step on with your bare feet? The water temperature will be lower than the asphalt. Water's *capacity* to hold heat is higher, meaning it takes much more heat energy to increase the water temperature compared to asphalt.

  Takeaway – Water cooling can effectively transfer heat from a processor using a significantly smaller volume of coolant compared to the air volume required for equivalent air cooling.

What does this mean in terms of energy costs? First, **when you cool your IT equipment with water, you can set your chilled water temperature higher compared to air**. For every 0.6°C (1°F) increase in chilled water temperature, the chiller efficiency increases 1%-2%.[8] For variable speed chillers, the efficiency increases 2%-4%. And this doesn't include the potential energy savings from economizer hours when the chiller is off or at partial load, which can be significant in cooler climates. We recommend you use the *Cooling Economizer Mode PUE Calculator* to see the impact of IT inlet air temperature on economizer hours for a given data center location. Secondly, distributing a high volume of air uses more energy (fans) than the pump energy used to distribute a smaller volume of water.

What does this mean in terms of capital costs? If you want to maximize the efficiency savings from increased chilled water temperatures, you will incur capital costs for a separate chiller plant. This leaves you with two basic choices:

1. Use the same "low-temperature" chiller for both loads and forego the energy savings offered by liquid-cooled IT.
2. Invest in a separate "high-temperature" chiller plant dedicated to liquid-cooled IT loads (higher chilled water temperature) and use the existing chiller for air-cooled IT loads.

GUIDANCE:
We recommend you perform a total cost of ownership (TCO) analysis to determine which approach will yield the lowest net present value (NPV), over the next 10 years or more. Whenever economizer hours are involved, the data center location plays a

[7] White Paper 265, *Liquid Cooling Technologies for Data Centers and Edge Applications*, page 12

[8] Arthur A. Bell, Jr., *HVAC Equations, Data, and Rules of Thumb*, (New York: McGraw-Hill, 2000), p. XX

material role in the TCO analysis. Expect the NPV difference to increase in extreme climates.

As a rule of thumb, the TCO analyses will favor different approaches based on whether the data center is located in a hot or cold climate:

- In **hotter** climates, the high-temperature chiller plant yields significantly more economizer hours compared to the low-temperature chiller plant. The energy savings will tend to **favor investing in a separate high-temperature chiller plant**.
- In **colder** climates, the high-temperature chiller plant yields *marginally* more economizer hours compared to the low-temperature chiller plant. The smaller the difference in economizer hours, the lower the energy savings, and the more difficult it is to justify the capital cost of a separate high-temperature chiller. The energy savings will **tend to favor using the same low-temperature chiller plant to support both air-cooled and liquid-cooled IT loads**.

When using the same chiller, a data center operator has a narrow band of chilled water temperature setpoints from which to choose. This becomes a constraint on economizer hours. With a single chiller plant supporting air- and liquid-cooled IT loads, the control logic becomes even more critical. It must measure, analyze, and then apply specific control logic (i.e., rules) to the whole system to maintain cooling while preventing condensation (dew point control).

Figure 4 illustrates the chilled water temperature relationship between air-cooled and liquid-cooled IT.[9] The dashed line represents the chilled water temperature entering the CDU as a function of the *IT coolant* temperature. The solid line represents the chilled water temperature entering the CRAH as a function of the *IT air* temperature. The yellow shading highlights the narrow band of chilled water temperatures you are limited to when using the same chiller for both liquid and air cooling. The IT temperatures on the x-axis, for liquid and air cooling, are based on ASHRAE's TC 9.9 Datacom Encyclopedia.

Figure 4

Chilled water temperature relationship between air-cooled and liquid-cooled IT

Assumes hot aisle containment (air cooling) and 4°C (7.2°F) CDU approach temperature (liquid cooling)

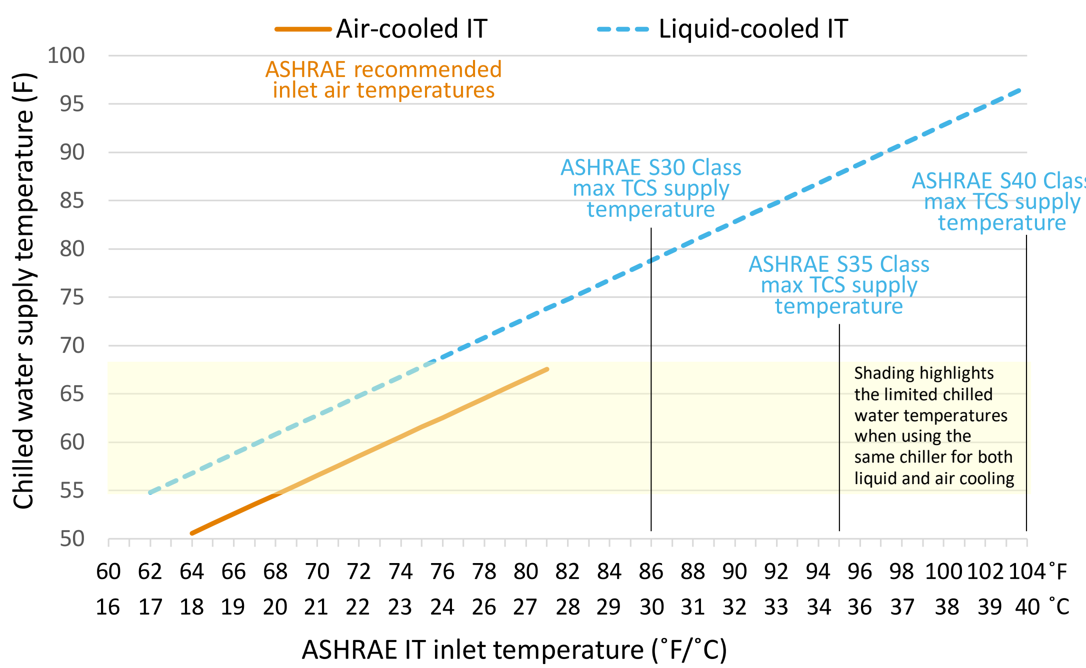

[9] Assumptions like location, % load, CDU approach temperature, and CRAH type will change the graph values, but the conclusion of a narrow band remains the same.

There are two additional considerations that may influence the selection between the two approaches:

- The continued upward trend of GPU power density means that TCS coolant temperatures will need to decrease. Therefore, in order to support multiple generations of IT, key industry players have proposed a minimum TCS coolant temperature of 30˚C (86˚F).[10]
- The concept of heat recovery may provide additional justification for investing in a separate high-temperature chiller plant. A heat recovery system uses the high-temperature return water (after going through the liquid-cooled IT) to assist in other heating applications like district heating, fish farms, etc.

3. Direct coupling between servers and DLC infrastructure complicates specification

With air cooling, there are no pipes or ducts directing air to or from an *individual* server. At most you could have air ducted to a single rack of servers using rack air containment. In essence, air-cooled servers and the air conditioning units are only *loosely* coupled. In contrast, in a DLC system, the coolant distribution infrastructure is physically connected (i.e., *directly* coupled) to the IT equipment. The challenges associated with direct coupling can be categorized into 3 distinct topics.

- Distribution piping
- Liquid used as the cooling fluid
- TCS must be designed to meet individual server cooling needs

Distribution piping is required for each individual server to connect the CDU to the server's cold plates. Some examples of how this complicates specification include:

- Piping is an additional server connection that requires mechanical compatibility between quick-connect couplings (e.g., between IT and manifold).
- Piping occupies more space, compared to thin electrical wires, making redundancy less feasible inside the rack and server.
- Flexible cooling pipes (i.e., hoses) shouldn't exceed their bend radius limit. Exceeding bend radius increases the risk of leaks and impedes coolant flow.
- Space must be allocated for hoses and manifolds without jeopardizing air flow for air-cooled components, even though IT racks are already space-constrained.
- Pressure drops vary across the TCS loop, potentially causing insufficient cooling to one or more servers.
- Pressure drops vary across different server models installed in the same rack and can cause insufficient cooling to one or more servers.

Liquid used as the cooling fluid instead of air means that specification must comprehend the risk of leaks inside racks and servers (as opposed to stopping at the CRAH). Specification is also required for a coolant filtration system, treatment system, and material compatibility to prevent corrosion.

TCS must be designed to meet individual server cooling needs**.** Liquid-cooled server components are completely dependent on centralized CDU pumps and controls for cooling. It's analogous to pushing on a string. The CDU must provide

[10] Mills et al., *30˚C Coolant - a Durable Roadmap for the Future*, June 2024

enough coolant flow through each server. In contrast, air-cooled servers control 100% of their own air flow.

GUIDANCE:
Liquid-cooled IT equipment and the TCS in general, have stringent requirements regarding fluid temperature, pressure, flow rate, material compatibility, and coolant quality. These requirements are provided by the IT vendor and help specify the TCS piping loop which includes the CDU. In essence, all this information starts to describe your liquid cooling system.[11] Figure 5 illustrates key design parameters and type of vendors involved in the specification process.

Figure 5

Facility-side and server-side requirements that help specify the TCS

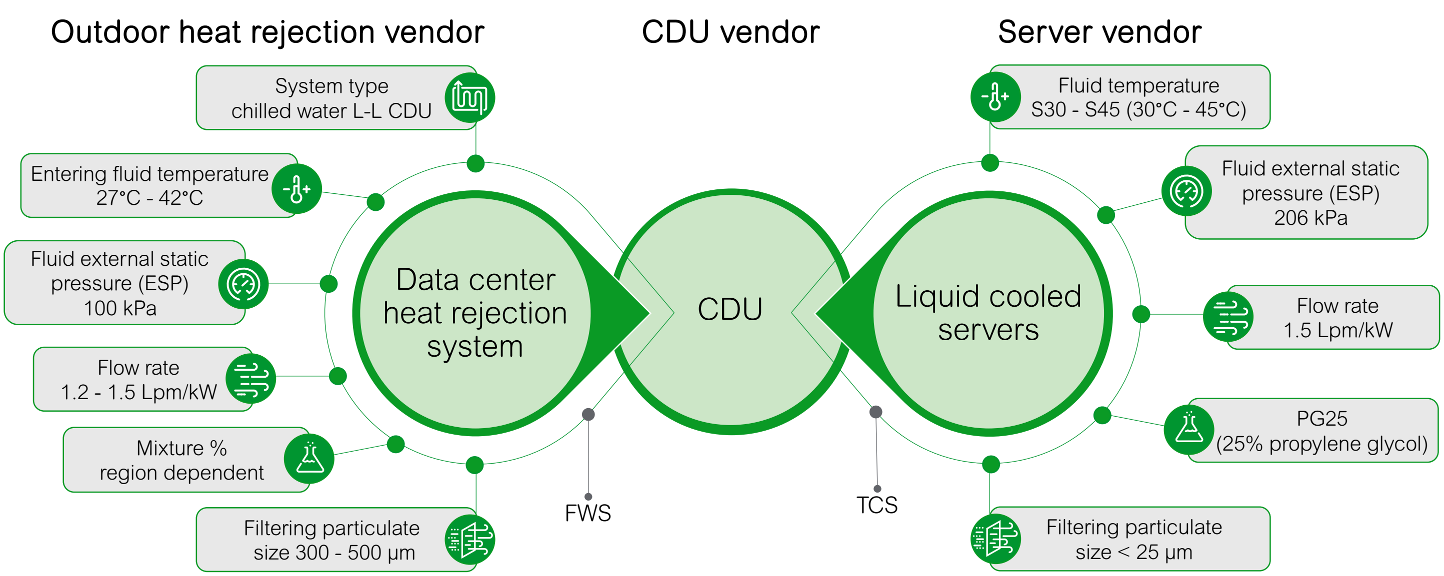

Schneider Electric DCRS, 2025

Below we provide best practice guidance on addressing each of the three attributes of direct coupling.

Distribution piping
Today there are no industry-accepted standards for quick-connect couplings, manifolds, piping, etc. to ensure that all these components fit together. Consequently, close collaboration with your IT vendors is essential for maintaining IT equipment warranties. It is also why we recommend a single TCS loop dedicated to a single server model. In other words, if you mix different server models in a TCS loop, there's a higher risk that some servers won't receive sufficient cooling.

Enterprises deploying AI clusters will have to wait for piping standards, which will likely come with help from server OEMs and/or hyperscalers deploying customized IT. In the meantime, enterprises should work with their server vendors who have partnered with liquid-cooling infrastructure vendors. Together they specify the server's quick connector couplings and rack manifolds so that they are compatible. In some cases, a rack-mounted L-to-L CDU is also specified for a given server. You'll want to ask for guidance on over-dimensioning piping and piping layout to avoid stagnant coolant which can breed contaminants.

[11] White Paper 133, *Navigating Liquid Cooling Architectures for Data Centers with AI Workloads*, describes DLC architectures and provides guidance on selecting the appropriate CDU.

Note that plastic piping has also been used in DLC systems. It can provide benefits over metallic piping including immunity to corrosion. Reliable plastic piping connections can be made through fusing (melting), akin to welding metallic piping.

Liquid used as the cooling fluid

Leaks tend to occur at couplings and joints, not at the piping walls. As best practice, we recommend specifying leak detection sensors for the piping most prone to leaks, including at connections (i.e., couplings, welds, threaded, crimped), in drip pans, by cooling coils, and by heat exchangers. In some cases, the server vendor includes leak detection within the server. We recommend confirming with your IT vendor if this is included or if it's an option.

Specifiers should include direct and indirect leak detection in the TCS loop (including the CDU):

- Direct leak detection includes cable leak detectors placed along the path of piping. Point detectors (also called spot sensors) are most effective when placed at the lowest point of a drip pan.
- Indirect leak detection infers a leak, not by the presence of water, but by the absence of it. For example, monitoring the water level of a buffer tank in a closed system will infer a leak if the level drops. Or if air is getting into a loop, it's likely a sign that there's a leak somewhere. A turbidity sensor measures the particles in a liquid by how much light is deflected off of the particle. Air bubbles are also detected as a particle. Note that, in general, these sensors don't distinguish between air bubbles and other particles. In either case, it's important that the facility team investigate the issue to determine if they have a filtering problem or a leak or both. Air bubbles can also contribute to contamination in the TCS loop.

Specifying lighting in areas with higher leak risks is also recommended to help reduce the impact of long-term leaks. Sometimes visual inspection is the best monitoring solution. Over time, leaks leave visual traces like minerals, which could be easily noticed with enough light. In areas like raised floor plenums, inside IT racks, in a sofit, etc., these traces are easy to miss. All these sensors must report to a centralized control system such as data center infrastructure management (DCIM) or a building management system (BMS). Stakeholders must decide how much of this control is given to automatic intervention (e.g., closing valves, shutting down servers) vs. notifications. If your organization is early in its DLC journey, we recommend notifications coupled with manual intervention. This approach will greatly simplify the control system. For more information on leak detection see page 22 of the Open Compute Project (OCP) paper, *Leak Detection and Intervention*.

Lastly, the coolant in the TCS loop must be filtered to remove contaminants. In addition to fouling the cold plate, debris can also jam the quick connects causing leaks. Most CDUs offer the option of either 50 µm or 25 µm filters. However, we also recommend consulting with your IT vendor to determine if you require a side-stream filter for smaller particles. This is sometimes required for cold plates with very small microchannels. See ASHRAE white paper (pages 25-30) and OCP guidelines (page 14) for more information on filtration.

TCS must be designed to meet individual server cooling needs

The TCS operation is only as good as the design. If a specifier fails to comprehend the circuit with the highest impedance, some servers in the loop won't be properly cooled. There are various control strategies to pump coolant through servers in a TCS loop, such as constant flow rate, constant pump speed, or constant differential pressure. Which one to use depends on an application's requirements. With AI or

HPC, the priority is to minimize the risk of overheating the GPUs. All other requirements are secondary, but specifiers can't ignore capital cost and energy cost.

In a production AI cluster, every "circuit" leading to an individual cold plate will have variations in impedance or resistance to flow. Variations are also introduced during maintenance operations such as replacing failed server nodes. The control logic must account for the variations in impedance without risking any server. An effective control strategy for this situation is constant differential pressure control.[12]

This method uses special "energy valves" to keep a constant pressure across the supply and return pipes of each rack. The energy valves are placed on the return side of each rack, and every valve setpoint must be configured based on the impedances of that circuit. A constant pressure means that the flow across the servers tends to remain steady even if a server is removed. This is good for the servers and it also saves energy by not over-pumping coolant through the TCS.

If you used this control method with different server models in the same loop, it wouldn't work at the rack level. The extreme (and very expensive) version of this method would be to place an energy valve at every server thereby allowing each server to have their own flow rate. While this would be ideal, especially for a mixed server loop, its cost doesn't outweigh the benefit, and it would be challenging to integrate all the valves into a production rack. Since there is no standard DLC control platform, we recommend specifiers leverage their infrastructure vendor for help with this control system.

Though we're in the early days of DLC technology, specifiers can leverage detailed documentation from industry organizations like OCP and ASHRAE who are helping to advance the adoption of DLC and provide specifiers helpful information for their DLC projects. Table 2 provides some of these useful documents:

Table 2

Industry documentation from OCP and ASHRAE to aide in specification of DLC projects

Industry organizations	Documentation
OCP	• *Hose and Manual Coupling – Best Practices* • *Universal Quick Disconnect (UQD) Specification* • *Guidelines for Using Propylene Glycol-Based Heat Transfer Fluids in Single-Phase Cold Plate-Based Liquid Cooled Racks* • *OCP OAI System Liquid Cooling Guidelines* • *30˚C Coolant - a Durable Roadmap For the Future* • *Leak Detection and Intervention*
ASHRAE	• *Water-Cooled Servers: Common Designs, Components, and Processes*, pages 34-43 cover couplings and flexible piping. • *Liquid Cooling Guidelines for Datacom Equipment Centers*, specifically chapter 4, Liquid Cooling Implementation for Datacom Equipment and chapter 6, Liquid Cooling Infrastructure Requirements for Technology Cooling Systems.

Reference designs are another tool to help specifiers with their projects. An example of this is Schneider Electric's *AI Reference Designs to Enable Adoption: A Collaboration Between Schneider Electric and NVIDIA*.

In the absence of standards, your ecosystem of partners becomes critical to reducing the risk of defects in your AI or HPC cluster. System integrators play a crucial role in this ecosystem by designing and implementing solutions that accommodate vendor specifications. Their expertise supports a final build that adheres to the vendor requirements, while maintaining vendor warranties.

[12] Shahi et al., *Experimental Study of Transient Hydraulic Characteristics for Liquid Cooled, Data Center Deployment*, Dec 2022

As an interesting aside, if we were talking about powering these server clusters, there would be no such challenges related to the "hard coupling" of wires, circuit breakers, transformers, etc. Why? Because global standards such as harmonized electrical codes and electrical connectors, answer these questions. These standards form a "common language" around powering these loads right down to the wire size. We believe that this too will become the case for liquid cooling as the industry matures.

4. Lack of CDU system efficiency standards complicates comparisons

Choosing between multiple CDUs can be confusing without a common rating and testing standard. Consider some of the variables influencing a CDU's efficiency or capacity rating: fluid types, flow rates, fluid temperatures, and ambient temperature. Without a standard, vendors are free to choose their own values for these variables and claim the best performance under different operating conditions. This makes it difficult to compare two or more CDUs, and challenging to specify the best CDU for your design.

For example, assume the rated capacity of "CDU A" is higher than the rated capacity of "CDU B". It's possible that "CDU A" provides a *lower* capacity when assessed at the conditions for your liquid cooling design. This isn't an issue with cooling components like CRAHs because they follow established testing and rating standards like, AHRI Standard 1361 (SI), *Performance Rating of Computer and Data Processing Room Air Conditioners*. A data center manager can confidently compare two or more CRAHs or chillers on the basis of efficiency and rated cooling capacity (kW). Like CDUs, this equipment has heat exchangers, motors, pumps, and fans. Yet today, despite being very similar, there are no such CDU standards.

GUIDANCE:
ANSI/ASHRAE Standard 127-2020, *Method of Testing for Rating Air-Conditioning Units Serving Data Center (DC) and Other Information Technology Equipment (ITE) Spaces*, "establishes a uniform set of test requirements for rating air conditioning units" applied to data centers. An upcoming addendum to this standard, *BSR/ASHRAE Addendum b*, will add DLC-related content including definitions, equations, test requirements, and test conditions for rating CDUs. Once released, vendors will have the necessary guidance to rate their CDU efficiency and capacity in a consistent manner. Until then, we recommend specifiers ask vendors to provide their ratings using the draft test points from Table 9-1 (page 14) in the draft version of Addendum b, along with the test procedures in section 9.

While having an apples-to-apples comparison between CDUs is beneficial, it doesn't define the efficiency of the entire DLC system, which is ultimately more important. The CDU, and the TCS loop it supports, must be specified along with the FWS. Together these systems determine the overall cooling efficiency *and* capacity. The biggest levers in determining the efficiency of a DLC system are:

- Data center climate – The colder the climate, the more hours the system can operate without using energy-intensive compressors.
- Server heat capture ratio – This is the percentage of heat captured by liquid and captured by air. For example, 80% heat capture by liquid means that 20% is captured by air. Because DLC systems are significantly more energy efficient than air-cooled systems, capturing more server heat with DLC lowers data center energy consumption.
- Evaporative cooling – Spraying water over the outdoor airstream cools the air through evaporative cooling. This cooler air lowers condenser water or

refrigerant temperatures and saves compressor energy. As expected, it is not recommended for water-restricted locations.

- FWS and TCS coolant temperatures – The higher the coolant temperatures in the FWS and TCS loops, the less energy is consumed in rejecting heat to the outdoors. These temperatures are limited by the maximum TCS temperatures allowed by the liquid-cooled chips and the CDU approach temperature (discussed below). 30˚C (86˚F) is the current industry maximum as discussed in a previous challenge.
- Chiller efficiency – In warmer climates, the chiller efficiency plays a larger role in energy consumption because more compressor hours are required.

Importance of CDU approach temperature

The approach temperature is the difference between the temperature of the coolant leaving the CDU (entering the IT servers), and the temperature of the facilities chilled water entering the CDU from the FWS loop. The lower the CDU approach temperature, the higher you can specify the chilled water temperature. Higher chilled water temperature increases economizer hours and reduces chiller compressor energy. State of the art CDU approach temperatures range from 2˚C (3.6˚F) to 4˚C (7.2˚F), although some vendors provide as high as 6˚C (10.8˚F).

While striving for efficiency is a good thing, keep in mind that switching to a DLC system is already a step change in efficiency. We recommend you weigh how much of your time is spent specifying for efficiency vs. for reliability. For example, it may be wise to spend more time ensuring that TCS sub-systems work properly (e.g., flow control, leak detection, coolant treatment, resiliency/redundancy).

A final note related to CDU capacity. For data center owners just starting their journey into DLC systems, we recommend specifying CDU capacities that account for less than 10-20% of your liquid-cooled servers. This will limit the disruption to data center operations if a failure occurs on a single TCS loop.

In other words, put as many servers into a single loop as your risk tolerance allows.

5. Provisioning IT space for unknown liquid-cooled IT risks higher costs from stranded capacity

In many cases data center owners must provision their IT space to support liquid-cooled IT without knowing what specific liquid-cooled IT equipment will show up. This scenario is especially common with colocation providers planning for new tenants. At face value this doesn't appear to be a major issue since this happens regularly with air-cooled IT spaces. But unlike air-cooling, DLC directly couples the coolant supply to every single server.

This means that every CDU, and its TCS loop, needs to be capable of supporting a maximum quantity of racks at the lowest density as well as the minimum number of racks at the highest density. The wider the range of this min and max, the greater the likelihood that you will strand some piping and valves by deploying racks at higher densities. There's also a higher expense for provisioning TCS distribution piping and valves. If you deploy the maximum number of racks (lowest density), you overpay for oversized piping. If you deploy the maximum rack density (lowest rack quantity), you overpay for stranded branch piping. As an example, imagine you have a 300 kW CDU supporting a range of 3 racks (100 kW/rack) to 10 racks (30 kW/rack). If you support 10 racks, you strand 70% of the pipe's and manifold's cooling capacity. If you support 3 racks, you strand 7 of the 10 distribution pipes feeding each rack. While it's true that you could order the manifolds *after* getting a committed customer, the piping tends to commit you to a specific density.

GUIDANCE:
Unfortunately, today there are no TCS solutions that offer the same flexibility of air-cooled systems. A good approach is to specify all the bulk cooling systems first, including chillers, pumps, cooling towers, dry coolers, and CDUs. The TCS main loop piping diameter is sized according to the CDU kW capacity and isn't influenced by the rack densities. However, the length of this piping should accommodate the maximum number of racks for the loop, in other words, the lowest rack density (kW/rack). The branch piping feeding the rack manifold in each rack must be sized for the maximum rack density. Like the branch piping, the manifold must be oversized to support the highest density rack. Rack manifolds can typically be ordered once you know the number of racks.

Another approach is to specify different power densities in different areas within the data center. This method does not eliminate the risk of stranding power or cooling infrastructure, but it does offer flexibility for supporting lower density deployments like inference workloads. Another suggestion is to develop reference designs for two or three DLC "pods" at different densities. Depending on the data center layout, these pods could be staged on skids with bulk cooling components pre-installed with CDU, filtration, coolant treatment, etc. When the IT equipment is confirmed, these pods could be finished off with the TCS loop design that is best matched to the IT.

Note that this challenge is similar to specifying power distribution to undefined IT equipment. Power distribution is also "coupled" to the IT equipment and therefore must be oversized to accommodate the highest rack density.

Installation

For the installation phase, we focus on one key challenge that data center stakeholders are facing. As adoption of DLC grows, we anticipate adding new and different installation challenges and guidance.

6. Complexity of preventing TCS contamination increases the risk of server damage

When installing the TCS loop, contamination is a bigger concern compared to the FWS loop because it poses serious consequences for the IT equipment. Contaminants can foul cold plates, leading to poor cooling performance and damage. Debris can also jam quick connects causing leaks. Contamination can be organic, such as living biological systems or their secretions; or inorganic, such as metal ions, salts, and debris or suspended particles. Any matter that is not meant to be in the coolant, constitutes contamination.

TCS fluid contamination comes through four primary mechanisms:

- Fouling – Material suspended in the coolant can accumulate at specific locations in the TCS loop causing clogging and increasing the corrosion potential at the accumulation sites. Fouling in heat exchangers (IT cold plates and CDUs) also reduces their efficiency. It becomes an additional heat resistance layer that must be overcome to transfer the heat from the IT load.
- Corrosion – Metal ions are removed from the wetted surfaces and can travel with the coolant to any location within the TCS loop. These ions can re-deposit at a different location causing clogging. As corrosion progresses, larger particles are removed from the wetted surfaces and travel at high-speed within the coolant. These particles cause abrasion of the surfaces they encounter, for example the microchannels in the cold plates, and can cause their premature failure. Corrosion occurs when metal surfaces are exposed to a liquid (the

coolant in this case). How fast or slow corrosion progresses depends on the specific metal and liquid acidity/alkalinity conditions, showing increased rates of corrosion when pH deviates from neutral.

- Microbial formation – Specific types of bacteria, fungi, and algae grow at stagnation points along the TCS loop, such as unused (future capacity) distribution feeders. These microbes can feed from metal surfaces, and/or secrete substances that cause the coolant pH to shift from its usual range (pH 8 - 9.5), to acidic (pH < 7), or alkaline (pH > 9.5) conditions, accelerating corrosion.
- Scale formation – Salts dissolved in fluid recrystallize over the TCS surfaces, creating sites for fouling. As scaling progresses and crystallized salts become larger, they could break free and accumulate at various TCS sites, or abrade TCS surfaces while traveling suspended in the coolant.

Potential sources of contamination during the installation of the TCS loop include:

- incompatible materials used for connectors, piping, manifolds, etc.
- the use of incompatible fluid sealants for treaded joints. Incompatible sealants can contaminate the cooling system by reacting with the coolant, releasing particles into the fluid, or altering coolant properties.
- debris produced during manufacturing of TCS components. Various types of debris can be generated, such as metal shavings from machining operations, plastic particles from molding or cutting processes, and dust or other fine particulates from assembly procedures.
- microbial formation within the coolant at TCS stagnation points. Depending on the environmental conditions, and precautions taken, if installation is interrupted for example, and TCS lines are left filled and stagnant for a couple of weeks, bacterial formation can occur.
- careless manipulation of connected components. Improper handling of TCS components during installation can lead to contamination. This includes not wearing gloves when handling parts, using inappropriate cleaning materials (e.g., non-lint-free wipes), and working in dusty or unclean environments.
- IT equipment installation in the TCS loop. When servers are installed at unused feeders of the TCS, stagnant coolant in unused lines may contain contaminants. Connecting servers without flushing can introduce these contaminants to server cold plates.

GUIDANCE:
Keeping the TCS free from contamination is an important yet complex installation task, given the potential sources of contamination listed above.

We recommend a modular approach, limiting the size of the TCS loop to minimize the consequences of contamination. Contamination in one of ten loops risks damaging a smaller number of servers compared to a single loop with ten times the servers. Similarly, having multiple smaller TCS loops rather than larger ones can prevent microbial formation at stagnant sections, such as unused feeder lines.

Detailed installation procedures can sometimes seem excessive and unnecessary. However, we recommend that you take all the precautions and follow the IT and CDU manufacturers' installation guidelines with extreme care. Not doing so can lead to contamination, where the source can be very difficult to identify, and/or expensive to correct. This may also jeopardize equipment warranties.

We recommend the following practices during installation to avoid contamination:

- Prioritize the preassembly of components under environments with minimal contamination including particles (dust) and/or bacteria.
- Use welding to minimize dissimilar material piping couplings in the TCS and FWS loops. However, keep in mind that the high heat applied to the junction may locally damage the anti-corrosion properties of the piping materials. Apply passivation treatments in areas worked through high heat, followed by thorough washing to prevent contamination issues.
- Inspect all soldered or brazed junctions for cleanliness, as these sites are prone to contamination from leftover soldering/brazing fluxes.
- Avoid fluid sealants that could release particles into the TCS loop.
- Add biocide additives to the coolant prior to the first filling of the system to help prevent bacterial growth.
- Flush the entire TCS loop thoroughly before connecting CDUs and IT equipment. During commissioning, keep records and samples of all procedures and environmental conditions. Although specific procedures are given by the IT and cooling equipment vendors, you should record key environmental conditions such as temperature, pressure, and relative humidity.
- Purge any air out of the components prior to their installation to help prevent bacterial growth, as well as disturbances in the coolant flow.
- Take precautions to maintain cleanliness during addition or exchange of IT equipment in the TCS loop. Prior to installation, perform additional flushing and purge air from the system.
- Avoid time gaps during installation that may allow stagnant coolant to breed bacteria. If coolant must be stagnant for more than a week, circulate and filter it within the TCS loop until work resumes.
- Strictly follow installation and commissioning procedures. Don't underestimate the importance of handling components properly to avoid contamination.

We also believe it's crucial to leverage your ecosystem of partners. Vendor collaboration will help reduce uncertainty in all installation tasks to keep the TCS loop free of contamination. Moreover, equipment manufacturers have control of fabrication procedures and materials involved. Therefore, they become the starting point to prevent contamination. Work with the IT equipment, CDU, and cooling solutions vendors to validate the procedures to install, operate, and maintain your TCS loop free from contamination. We also recommend you ask vendors to ship components in antimicrobial packaging or inert atmospheres as needed; or filled with appropriate gas or fluids to prevent contamination during transit. For example, shipping components under nitrogen atmospheres prevents corrosion by displacing oxygen.

Operation

The most challenging portion of operating a DLC system is the TCS. Facility teams are comfortable with operating chilled water plants because they've done it for decades. But when it comes to distributing coolant to an IT rack, it's very new for most, and it's a steep learning curve. This is another phase where early adopters should depend on their ecosystem partners. Here we focus on two key challenges we believe data center operators will face. We anticipate adding new operations challenges as DLC systems evolve over time.

7. Lack of clarity between cooling and server vendor warranties creates undue stress

There are two key reasons this challenge exists. First, DLC is new to the data center industry resulting in a steep learning curve. Second, the physical (direct) coupling between servers and coolant distribution complicates the delineation of responsibilities between vendors. As a result, the warranties for some components and services of a DLC system may be ambiguous and/or contradict one another. This often leads to finger pointing between vendors and creates stress for end users. These situations may also delay repairs and lead to additional expenses. The following questions help to convey the ambiguous roles and responsibilities:

- If a leak occurs inside a liquid-cooled server, who is responsible… the server vendor, the CDU vendor, the fluid vendor, or another vendor?
- If a CDU filter maintenance procedure conflicts with a server vendor's recommendation, what should you do to prevent voiding either warranty or both?
- If a contractor introduces a defect during a DLC installation causing a server failure, who is responsible?
- If a prolonged material compatibility issue causes server damage, whose warranty applies?
- If a CDU supports servers from multiple vendors, with different fluid pressure requirements, how do you avoid warranty issues with any one server vendor?
- If there are two mating components with conflicting torque specifications, does adhering to one automatically void the other?
- If "degraded" coolant causes fouling in a cold plate, how do the coolant and server warranties specify a degraded fluid? Are they consistent?

GUIDANCE:
If we look back to the days when chilled water plants were first introduced to data centers, we'll find similar questions to those above. Like chilled water plants, much of the angst stemming from these questions will be alleviated with the introduction of industry standards related to direct liquid cooling. Further relief will come through collaboration between cooling vendors and IT vendors. The work of standards development has begun and there are already several documents available that you may find helpful in Table 2.

As standards publish and the ecosystem of vendors, integrators, and end users gain more experience, liquid cooling will become more predictable (e.g., lower defects, consistent requirements, less surprises). Until then, we provide the following guidance to mitigate warranty issues.

- Ask your IT vendors to agree on a common TCS supply temperature, flow rate, and pressure. Set your FWS to enable this TCS supply temperature by adjusting its temperature and or flow rate.
- Segregate IT equipment in vendor-specific TCS loops with uniform cooling requirements to avoid disputes regarding warranties and responsibilities between different IT vendors.
- Clarify responsibilities between the DLC system installer, DLC vendor(s), and IT vendors for the operation and maintenance of the system.
- Examine warranties for all TCS components and ask clarifying questions early on, as part of the specification process.
- Review TCS maintenance procedures with all vendors to avoid any steps that would jeopardize your operation or warranties. Make sure service manuals include procedures like regular filter replacement, coolant monitoring, and assigning an "owner" for the monitoring and maintenance of the coolant health. Clarify with the server vendor what coolant properties to maintain.
- Also periodically circulate stagnant coolant in redundant or backup loops to avoid bacterial growth that can foul cold plates.

- Maintain and filter the FWS chilled water loop, as this improves the CDU heat exchanger performance and lifespan.
- Perform a thorough flow distribution analysis of the rack coolant distribution with actual server impedance data. This is similar to CFD analysis in air cooled systems to avoid hotspots.

8. Slow DLC system response to GPU power transients poses a risk of GPUs overheating

The chips that perform the mathematical operations to train models are collectively called "accelerators" but the most common of these are GPUs. During training, the power consumption of all GPUs rise and fall simultaneously. About four times per second, GPUs may peak to 150% or more of their thermal design power (TDP).[13] These peaks, which we call power transients, present a challenge to cooling operations. These peaks can last for up to 50 milliseconds and increase the risk of overheating the GPUs. If the GPUs overheat, it's possible that the server throttles (slows down) or shuts down the GPUs. In the worst case, one or more of these expensive GPUs could be damaged.

Is it likely that transient peaks will shut down or damage GPUs? In any electrical circuit, electrical energy is converted into heat energy. While electrical energy propagates through a circuit at nearly the speed of light, heat energy moves much slower. How fast the heat energy propagates from the GPU die to the cold plate depends on many variables including the temperature difference between the two. If we let the GPU operate at 100% utilization (i.e., the chip's TDP), then all the components would eventually reach a steady state temperature. Using a steady state heat transfer analysis,[14] we estimated the GPU temperature to be 78˚C. We then modeled the GPU utilization at 150% of TDP (1,050 W) for 50 ms, and at 100% (700 W) for 200 ms. We modeled this for 1 second and plotted the temperatures and power consumption in Figure 6.

Figure 6

GPU temperature modeled at various power consumption ranging from 700W to 1050W at various TCS temperatures

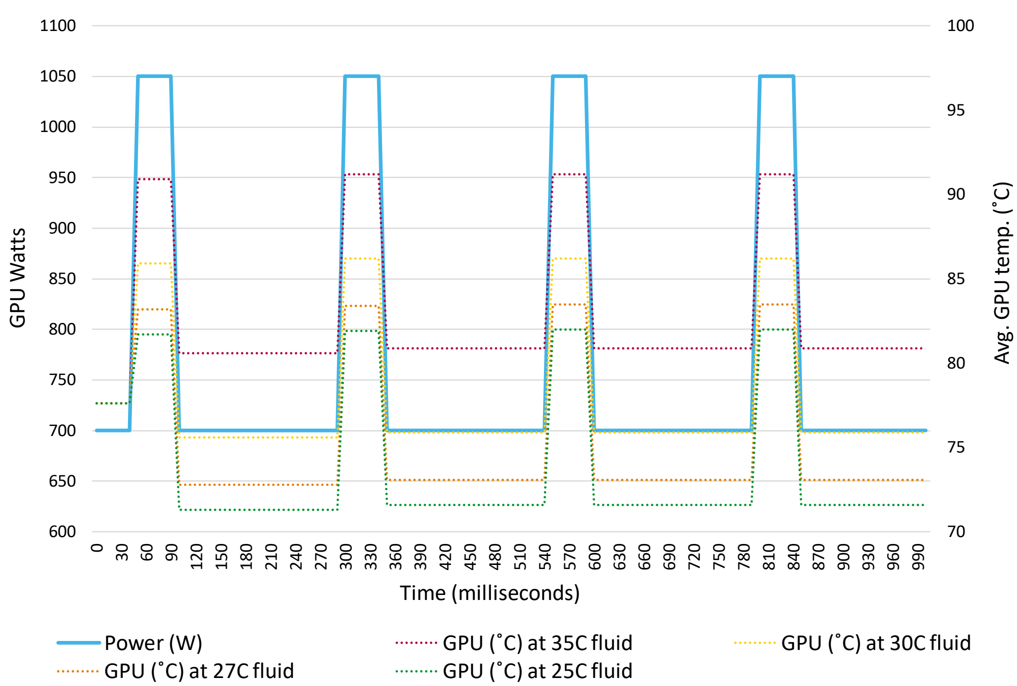

[13] Daniel Bizo, *Erratic power profiles of AI clusters: the root causes*, Uptime Institute, Sept 2024

[14] Major assumptions include copper cold plate area 2400mm^2, 10mm high, and convective heat transfer coefficient h=25,060 W/m²K. Cooling liquid (25% glycol) at 1 l/min supplied to the cold plate at 35˚C. "GPU temperature" is an averaged temperature of the entire GPU die volume.

In this analysis, we estimated that the GPU temperature increased approximately 10˚C between transients to a high of 91˚C. This temperature tends to be on the higher end of GPU limits. While this analysis doesn't prove that GPU power transients will or won't cause thermal issues, it does suggest a real possibility.

GUIDANCE:
Today there is no conclusive and publicly available evidence to confirm that GPU power transients lead to overheating issues. It's possible that certain GPU models and package configurations are better than others in this regard. The uncertainty is made worse by the lack of standard DLC system designs and practices. What may work for one system may not work for another. As the industry studies this challenge further, we provide the following guidance:

- Use software to power cap the GPUs in an AI cluster. This will certainly prevent GPUs peaking above their TDP. Unfortunately, this solution comes at the expense of increased training times, something many organizations reject.
- Set up a test plan to monitor your IT and cooling system. This includes GPU peak power, average power, CDU temperatures, pressure, and flow rate. After commissioning is complete, run test workloads on servers operating off a single TCS loop. It's important to monitor GPU temperatures across all servers, even if they are all the same models. There could be inconsistent cooling flow among the servers in a loop which could starve a particular server thereby increasing GPU temperatures. Start with a low % GPU utilization, and as you increase the utilization, note the trend in GPU temperatures.

  If any GPU temperature[14] is likely to reach 90˚C, decrease the cold plate's fluid temperature. The following example is for illustrative purposes only. If the average GPU temperature is 90˚C and the cold plate average temperature is 40˚C, the temperature difference (deltaT) is 50˚C. Reducing the cold plate to 35˚C increases the deltaT to 55˚C which increases the rate of heat transfer between the GPU and cold plate. The most effective way of reducing the cold plate temperature is to reduce the chilled water supply by 5˚C. Doing so keeps the flow rates the same on both the TCS and FWS loops. See Figure 6 for GPU temperatures with different TCS/cold plate temperatures.
- Increase the TCS flow rate to lower the average temperature across the cold plate. Note, this isn't as effective as reducing the TSC fluid temperature into the cold plate as discussed above. Decreasing the TSC or FWS temperature or decreasing their deltaT will decrease the DLC system energy efficiency. This may be an acceptable trade off in the early phases of your AI deployment.
- Save your system baseline data to compare later in case of anomalies.
- Evaluate the latest server (e.g., GPU) firmware, drivers, and updates. Vendors sometimes release updates that may improve thermal or power management.
- Continue to monitor and analyze your GPU and system temperatures for the first few weeks or months of operation to provide the confidence that the system is running as it should.
- Stay informed on industry news regarding this issue as there are numerous stakeholders investigating it, including your vendors.

Next steps

This paper covered three main data center life phases: Specification, Installation, and Operation. Given that many data center professionals are just starting their direct liquid cooling journey, we recommend four key next steps to help you successfully specify your liquid cooling architecture.

- **Identify a design firm, cooling vendor, and integrator** to help you design your DLC system. Ask for existing **reference designs** as a starting point for your own DLC system.
- **Collaborate with your IT vendor(s)** to understand your IT needs. This is the single most important factor in designing and specifying a reliable, cost effective, and efficient DLC. Once you've selected a liquid-cooled server model, you will know the recommended coolant temperature, flow, pressure, and chemistry. You should also know what the server's maximum power consumption, heat capture ratio, coolant deltaT across the server, quick connect specification, coolant filter requirements, and other information required by your TCS vendor(s).
- Once you've selected your cooling and IT equipment, **collect and analyze warranty terms and conditions** to identify contradictions or conflicts. Clarify any questions that arise with the associated vendor(s).
- **Build a list of common wetted materials**. If you don't have one started, begin with those from *ASHRAE* (Table 2, page 28) or *OCP* (Table 2, page 11). Try to narrow down this list by asking various server vendors which materials are NOT compatible with their cold plates. This should serve to reduce your list to a few materials common among all the server vendors. The shorter the list, the fewer metals become part of your TCS, and the lower the likelihood of leaks and server cooling issues.

About the authors

Victor Avelar is a seasoned expert in data center energy efficiency and design, serving as the Chief Research Analyst in Schneider Electric's Data Center Research & Strategy group. With over 25 years of experience, Victor leads cutting-edge research and best practice development for sustainability, risk management, and next-generation data center technologies. He's a trusted advisor to clients globally, providing actionable insights on enhancing infrastructure performance through innovative solutions such as liquid cooling and energy modeling. Known for his clear, practical guidance, Victor helps organizations tackle the evolving challenges of sustainable and efficient data center operations. He is central to the development of technology adoption forecasts for data centers. He also leads the peer review process for all DCRS content. Victor holds a bachelor's degree in mechanical engineering from Rensselaer Polytechnic Institute and an MBA from Babson College. He is a member of AFCOM and a sought-after speaker on AI infrastructure.

Dr. Maria A. Torres Arango is a Research Analyst in Schneider Electric's Data Center Research & Strategy group. She investigates technology, materials and infrastructure to guide data center strategies. Her current work focuses on energy storage and liquid cooling technologies. She also examines market forces driving technology advancement and participates in developing tools that highlight tradeoffs or selection in critical data center decisions. Maria's former expertise involves materials design and optimization; and fundamental studies on materials synthesis processes using X-ray characterization at the National Synchrotron Light Source II, Brookhaven National Laboratory. A lifelong learner, Maria holds a BS in aeronautical engineering from Universidad Pontificia Bolivariana, Colombia; and a MSc in aerospace engineering and a PhD in materials science and engineering from West Virginia University.

Stefano Lena has been working in the HVAC industry since 2006. After earning his master's degree in electronic engineering, he joined an Italian HVAC manufacturer as a product development engineer. Stefano next moved to an HVAC sales and key account management position for a Swiss HVAC company. In 2017, he joined Schneider Electric as product manager and worked in marketing roles developing his expertise in data center applications. In 2020 he joined the new Innovation and Strategy Cooling Team where he's been supporting cooling innovation projects, most importantly the liquid cooling program. Currently he's one of the subject matter experts on liquid cooling systems. His research is focused on liquid cooling system flow management.

David McGlocklin has been working as an engineer in the IT and data center cooling business for 25 years. Spending most of his career in sustaining engineering, he most recently led the new product design & lab teams in St. Louis before joining the cooling innovation & strategy team. He received his bachelor's degree in mechanical engineering from Georgia Tech. Prior to joining APC/Schneider Electric, David worked as an engineer at Carrier Transicold developing transport refrigeration. He is very active in ASHRAE and AHRI, serving on many committees and is the current chair of SSPC 127.

RATE THIS PAPER ★★★★★

Resources

AI Reference Designs to Enable Adoption: A Collaboration Between Schneider Electric and NVIDIA
Executive Brief 1

Optimizing AI Infrastructure: The Critical Role of Liquid Cooling
Executive Brief 3

The AI Disruption: Challenges and Guidance for Data Center Design
White Paper 110

Navigating Liquid Cooling Architectures for AI Workloads
White Paper 133

Liquid Cooling Technologies for Data Centers and Edge Applications
White Paper 265

Water-Cooled Servers: Common Designs, Components, and Processes
ASHRAE

Thermal Guidelines for Data Processing Environments
ASHRAE

Liquid Cooling Guidelines for Datacom Equipment Centers
ASHRAE

Leak Detection and Intervention
OCP

Hose and Manual Couplings – Best Practices
OCP

Universal Quick Disconnect (UQD) Specification
OCP

Guidelines for Using Propylene Glycol-Based Heat Transfer Fluids in Single-Phase Cold Plate-Based Liquid Cooled Racks
OCP

OCP OAI System Liquid Cooling Guidelines
OCP

30°C Coolant – A Durable Roadmap for the Future
OCP

Browse all white papers
whitepapers.apc.com

Cooling Economizer Mode PUE Calculator
TradeOff Tool 11

Browse all TradeOff Tools™
tools.apc.com

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Data Center Research & Strategy
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm